arXiv:1404.6287v1 [cs.DS] 24 Apr 2014

Improved Approximation Algorithms for Earth-Mover Distance in Data Streams

Arman Yousefi and Rafail Ostrovsky

University of California Los Angeles
{armany, rafail}@cs.ucla.edu

Abstract. For two multisets S and T of points in $[\Delta]^2$, such that $|S| = |T| = n$, the *earth-mover distance* (EMD) between S and T is the minimum cost of a perfect bipartite matching with edges between points in S and T, i.e., $EMD(S,T) = \min_{\pi:S\to T} \sum_{a\in S} ||a-\pi(a)||_1$, where π ranges over all one-to-one mappings. The sketching complexity of approximating earth-mover distance in the two-dimensional grid is mentioned as one of the open problems in [16, 11]. We give two algorithms for computing EMD between two multi-sets when the number of distinct points in one set is a small value $k = \log^{O(1)}(\Delta n)$. Our first algorithm gives a $(1+\epsilon)$-approximation using $O(k\epsilon^{-2}\log^4 n)$ space and works only in the insertion-only model. The second algorithm gives a $O(\min(k^3, \log \Delta))$-approximation using $O(\log^3 \Delta \cdot \log\log \Delta \cdot \log n)$-space in the turnstile model.

1 Introduction

For a metric space X endowed with distance function d_X, the *earth-mover distance* (EMD) between two multisets $S, T \subseteq X$, where $|S| = |T| = n$ is defined as $EMD_X(S,T) = \min_{\pi:S\to T} \sum_{a\in S} d_X(a, \pi(a))$ where π ranges over all bijections $\pi : A \to B$. In this paper we mostly deal with earth-mover distance over ℓ_1. Thus, when the metric X is ℓ_1 we omit the subscript and write $EMD(S,T) = EMD_{\ell_1}(S,T) = \min_{\pi:S\to T} \sum_{a\in S} ||a-\pi(a)||_1$. Earth-mover distance over the two dimensional plane has received significant interest in computer vision because it is a natural measure of similarity between images [25, 24, 22, 12]. Each image can be viewed as a set of features and the distance is the optimal way to match various features of images, where the cost of such a matching corresponds to the sum of distances between the features that were matched. Apart from being a popular distance measure in graphics and vision, variants of earth-mover distance known as *transportation cost* are used as LP relaxations for classification problems such as 0-extensions and metric labeling [2, 5, 6].

In this paper we study two-dimensional earth-mover distance in the streaming scenario. In the streaming scenario for earth-mover distance, the two multisets of input points are revealed to the algorithm as a stream of labeled points. The algorithm maintains a short "sketch" of the data that can later be used to estimate the cost of the optimal matching. Note that the algorithm does not produce an approximately optimal matching but only estimates its cost. The

stream can be viewed as a sequence of m operations where each operation either adds a point to one of the two multisets or removes a point from one of them. The streaming model in which insertion and deletion of points are both allowed is referred to in the literature as the *dynamic data stream* model. It's also called the *turnstile* model. The alternative is a streaming model in which only insertions of points are allowed, and such a data stream is referred to as *insertion-only* model.

Discrete geometric space In data stream scenario, we assume that points live in the discrete space $\{1, \ldots, \Delta\}^2$ (denoted by $[\Delta]^2$) instead of the continuous two-dimensional interval $[0, \Delta]^2$ where Δ is an integer upper bound on the diameter of the point set. This is not a common assumption in computational geometry where input points commonly have real coordinates. However, in real-life computations and in data stream algorithms, the discrete space is a common assumption because the input is assumed to have a finite precision.

Note that the assumption that the points of the two input multisets S and T live in the discrete space $[\Delta]^2$ implies that the distance between a point in S and a point in T is at least one. We assume that S and T are multisets, so multiple points of S (or multiple points of T) can share a location on the plane. However, we assume that no point of S shares location with a point of T.

1.1 Previous Results.

Computing the earth-mover distance is a fundamental geometric problem, and there has been extensive body of work focused on designing efficient algorithms for this problem [14, 23, 6, 12, 18, 19]. The challenge in designing efficient streaming algorithms for earth-mover distance is to construct and maintain a small space representation (or sketch) of both multisets from which earth-mover distance between them can be approximated. In one dimension, the EMD between two multisets can be reduced to calculating ℓ_1 difference between two vectors representing the point sets in $[\Delta]$. If the number of points in each multiset is n, the ℓ_1 difference between two vectors of size Δ can be approximated within a factor of $1+\epsilon$ for any $\epsilon > 0$ using $O(1/\epsilon^2 \log n \log \Delta)$ (see Fact 2). Thus, the EMD between two multiset of points in one dimensional space over a dynamic data stream can be approximated within a factor of $1 + \epsilon$ using $O(1/\epsilon^2 \log n \log \Delta)$ space. This is a folklore result, and the interested reader is referred to [3] for a detailed explanation.

In [8], Indyk gives a $O(\log \Delta)$-approximation algorithm for estimating the EMD between two multisets in $[\Delta]^2$ in one pass over the data that uses $O(\log^{O(1)}(\Delta n))$ space. His algorithm uses a probabilistic embedding of the EMD into ℓ_1 that has $O(\log \Delta)$ distortion [12, 6]. Later, Naor and Schechtman [17] showed that any embedding of EMD into ℓ_1 must incur a distortion of at least $\Omega(\sqrt{\log \Delta})$, so it is not possible to approximate EMD over a data stream within a factor better than $\Omega(\sqrt{\log \Delta})$ by embedding EMD into ℓ_1.

In [4] Andoni et al. gave a $O(1/c)$-approximation algorithm for estimating EMD in the two-dimensional grid $[\Delta]^2$ using space $O(\Delta^c \log^{O(1)}(\Delta n))$ for any

$0 < c < 1$. Their algorithm uses the result of [10] which decomposes the cost EMD over $[\Delta]^2$ into a sum of closely related metrics called EEMD, defined over $[\Delta^\epsilon]^2$. Each component of the sum is a sub-matching between subsets of the two original multisets. In [10] Indyk shows how to estimate the sum of sub-matchings by sampling sub-matchings using a random distribution where the probability of choosing a sub-matching is roughly proportional to its cost. In [4] the authors show how to approximate the sum of sub-matchings over a data stream.

For earth-mover distance in high dimensions, Khot and Naor [21] show that any embedding of EMD over the d-dimensional Hamming cube into ℓ_1 must incur a distortion $\Omega(d)$, thus practically losing all distance information. Andoni et al. [1] circumvent this roadblock by focusing on sets with cardinalities upper-bounded by a parameter s, and achieve a distortion of only $O(\log s \cdot \log(d\Delta))$. As a result, they show a $O(\log s \log d\Delta)$-approximation streaming algorithm that uses $O(d \log^{O(1)}(s\Delta))$ space.

1.2 Our Results

In this paper we give two streaming algorithms for approximating EMD in the two-dimensional grid $[\Delta]^2$ when the number of distinct points in one of the multisets is polylogarithmic. This is an interesting case because in applications the feature sets of images usually have bounded size. A similar case for high dimensions has been studied before in [1], but our constraint on the input is more relaxed than that of [1] because we require a bound on the number of distinct points in only one of the multisets while in [1] they require a bound on the size of both sets. The special case of EMD that we study is also important because of its connections to the *capacitated k-median problem* with hard constraints as we explain shortly.

Our first algorithm gives a $(1+\epsilon)$-approximation for any $\epsilon > 0$ using space $O(k\epsilon^{-2}\log^4 n)$. This algorithm uses coresets for k-median problem and it works in the insertion-only model. Our second algorithm works in the turnstile model (or dynamic geometric streams) and it gives a weaker approximation of $O(\min(k^3, \log \Delta))$ using $O(\log^3 \Delta \cdot \log\log \Delta \cdot \log n)$ space. Both algorithms naturally extend to work for higher dimensions. However, the second algorithm is better suited for higher dimensions because its memory usage does not depend exponentially on dimension d. The following table summarizes our results.

Algorithm	Approximation	space	Model
Algorithm 1	$1+\epsilon$	$O(k\epsilon^{-2}\log^4 n)$	insertion-only
Algorithm 2	$O(\min(k^3, \log \Delta))$	$O(\log^3 \Delta \cdot \log\log \Delta \cdot \log n)$	turnstile

Connections to Capacitated k-median Clustering The non-streaming version of Capacitated k-median clustering has been studied before (for example

[13, 15]), and it is known to be harder than k-median clustering with no capacities. In capacitated k-median clustering with uniform capacities over a data stream, in addition to a parameter k and a point set $P \subseteq [\Delta]^2$, we are given a parameter $c \geq n/k$. The goal is to find a set $Q \subseteq [\Delta]^2$ of size k that minimizes $\sum_{p \in P} ||p - f(p)||_2$ where $f(p)$ is one of the k centers that p is assigned to, and that the number of points assigned to each of the k centers doesn't exceed its capacity c.

Our algorithms for earth-mover distance can be extended to algorithms for capacitated k-median clustering with hard constraints. The input point set of the capacitated k-median clustering can be viewed as one of the point sets in the earth-mover distance and any set of k centers whose capacities add up to n can be viewed as the other multiset of points. The k-median cost of a point set respect to a given set of centers is the earth-mover distance between the input point set and the centers. The streaming algorithm for earth-mover distance can be used to keep a sketch of the input point set. At the end of the stream, the algorithm exhaustively searches all possibilities for k center points and, for each choice of k centers, all possible capacities of centers that do not violate capacity constraints and add up to n. For each possibility the algorithm approximates the earth-mover distance between the input point set and the capacitated centers and reports the centers with minimum value. Thus, the algorithm exhaustively searches all $\Delta^{O(k)}$ possibilities using small space and returns an approximate solution to the capacitated k-median problem with hard constraints. Note that the above algorithm does not violate capacity constraints. Thus, any of the algorithms in this paper can be turned into a streaming algorithm for the capacitated k-median clustering with hard constraints.

2 First Algorithm

In this section we show how to use coresets for k-median to give a $(1+\epsilon)$-approximation algorithm for EMD.

For a point set C and a point p, both in $\mathbb{R}^d$, let $d(p, C) = \min_{c \in C} ||p - c||_2$ denote the distance of p from C. For a weighted point set $P \subseteq \mathbb{R}^d$, with an associated weight function $w : P \to \mathbb{Z}^+$ and any point set C of k points, we define $Median(P, C) = \sum_{p \in P} w(p) d(p, C)$ as the the *price* of k-median clustering provided by C. In the k-median problem, the goal is to find a set C of at most k points in $\mathbb{R}^d$ such that $Median(P, C)$ is minimized. We also use $Median_{opt}(P, k) = \min_{C \subseteq \mathbb{R}^d, |C| = k} Median(P, C)$ to denote the price of the optimal k-median clustering for P.

Definition 1 (Coreset). *For a weighted point set $P \subseteq \mathbb{R}^d$, a weighted set $P_{core} \subseteq \mathbb{R}^d$ is a (k, ϵ)-coreset for the k-median problem if for every set C of k centers: $(1-\epsilon) \cdot Median(P, C) \leq Median(P_{core}, C) \leq (1+\epsilon) \cdot Median(P, C)$.*

Har-Peled and Mazumdar [20] prove the existence of small coresets for the k-median problem and show how to construct them. They also show how to construct and maintain coresets over data streams using polylogarithmic space when the points are only inserted into the stream. We use the following fact from [20].

Fact 1 (Theorem 7.2 from [20]). *Given an insertion-only stream P of n points in $\mathbb{R}^d$ and $\epsilon > 0$, one can maintain a (k, ϵ)-coreset of size $O(k\epsilon^{-d} \log^2 n)$ for k-median. The space used by the algorithm is $O(k\epsilon^{-d} \log^{2d+2} n)$ and the amortized update time is $O(\log^2(k/\epsilon) + k^5)$.*

2.1 Algorithm Description

Let $S, T \subseteq [\Delta]^2$ be two multisets of points such that $|S| = |T| = n$, where the number of distinct points in one of the sets is at most $k = \log^{O(1)}(\Delta n)$. Assume without loss of generality that the number of distinct points in T is k. The points of S and T are revealed to the algorithm in an insertion-only stream, the following algorithm computes an estimate of $EMD(S, T)$ as follows. The algorithm maintains a (k, ϵ)-coreset for k-median for the set S using Fact 1. For $d = 2$ the space needed to maintain the coreset is $O(k\epsilon^{-2} \log^4 n)$. Let S_{core} denote the coreset of S. The algorithm also keeps the entire set T of points in its memory. At the end of the stream the algorithm computes $EMD(S_{core}, T)$ using the "Hungarian" method [14].

We claim that $EMD(S_{core}, T)$ is a $(1 + \epsilon)$-approximation of $EMD(S, T)$. An important property of the coresets constructed for k-median in [20] that allows us to extend the use of coresets to earth-mover distance is the following. There is a one-to-one correspondence between the points of S and S_{core}. If for any $p \in S$, p' denotes the image of p in S_{core}, then the coreset construction guarantees that $\sum_{p \in S} ||p - p'||_2 \leq \epsilon \cdot Median_{opt}(S, k)$. Intuitively this means that each point of S is snapped to a point of S_{core} such that the sum of movements of points of S is at most $\epsilon \cdot Median_{opt}(S, k)$. It's easy to see from this property that for any set C of points $|Median(S, C) - Median(S_{core}, C)|$ is at most $\sum_{p \in S} ||p - p'||_2 \leq \epsilon \cdot Median_{opt}(S, k)$.

We now show how this property of the (k, ϵ)-coreset for k-median can be used to bound $|EMD(S_{core}, T) - EMD(S, T)|$. If for every point $p \in S$, its image in S_{core} is denoted by p', we have:

$$\begin{aligned} EMD(S_{core}, T) - EMD(S, T) &\leq \sum_{p \in S} ||p - p'||_1 \leq \sqrt{2} \cdot \sum_{p \in S} ||p - p'||_2 \\ &\leq \sqrt{2} \cdot \epsilon \cdot Median_{opt}(S, k) \leq O(\epsilon) \cdot EMD(S, T). \end{aligned}$$

The last inequality above holds because $Median_{opt}(S, k) \leq \sum_{p \in S} d(p, T) \leq EMD_{\ell_2}(S, T) \leq EMD_{\ell_1}(S, T)$. Thus we have shown that $EMD(S_{core}, T) \leq (1 + O(\epsilon)) \cdot EMD(S, T)$. We can also show that $EMD(S_{core}, T) \geq (1 - O(\epsilon)) \cdot EMD(S, T)$ using a similar argument. Thus we have:

Theorem 1. *For any $\epsilon > 0$ and any two multisets $S, T \subseteq [\Delta]^2$, where $|S| = |T| = n$, the number of distinct points in one set is bounded by k, and the points are revealed to the algorithm in an insertion-only stream, there is a one-pass streaming algorithm that approximates $EMD(S, T)$ within a factor of $(1 \pm \epsilon)$ and uses space $O(k\epsilon^{-2} \log^4 n)$.*

Why do coresets for dynamic data streams fail for EMD? A natural question to ask is if coresets can also be used for dynamic data streams where insertion and deletions are both allowed. Frahling and Sohler [7] proposed a method for constructing coresets that work for dynamic data streams. Their coreset construction is based on sampling points from the data stream, and it works for k-median, but it cannot be used for earth-mover distance. For a given coreset S_{core} of S, their algorithm constructs a set S'_{core} such that the point locations in S_{core} and S'_{core} are the same, but the weight of every point in S'_{core} differs from the corresponding point in S_{core} by a factor of at most $(1 \pm \epsilon)$. Thus, for any set C of k points, $Median(S_{core}, C)$ and $Median(S'_{core}, C)$ differ by at most a factor of $(1 \pm \epsilon)$ and computing the k-median cost for S'_{core} approximates the k-median cost for S_{core} and S. However, this argument does not work for earth-mover distance because $EMD(S_{core}, T)$ and $EMD(S'_{core}, T)$ may differ significantly. We mention a simple example to show that EMD is very sensitive to the weight of points in S_{core}, let T be a multiset containing two distinct points far from each other, each with weight $n/2$. Let also S_{core} be a coreset for S that contains exactly two weighted points, each with weight $n/2$. Each point of S_{core} is at distance one to a point in T. In this case $EMD(S_{core}, T) = n$, but changing the weights of points in S_{core} by a factor of $(1 \pm \epsilon)$ may affect the cost of $EMD(S_{core}, T)$ significantly.

3 Second Algorithm

Our first algorithm gives a $(1+\epsilon)$-approximation, but it doesn't work in dynamic geometric streams, and its space requirement is $O(k\epsilon^{-2}\log^4 n)$. We next present our second algorithm that works on dynamic geometric streams (when deletions are also allowed) and requires much less space specially for higher dimensions, but these advantages come at the cost of a weaker approximation ratio.

We start this section with some preliminaries and notations used in our description of the second algorithm and its related proofs. We use E^* to denote the set of edges of minimum-cost bipartite perfect matching between points of the two input multisets S and T. For an edge e that matches a point $p \in S$ with $q \in T$, let $||e||_1$ denote the ℓ_1 distance between e's endpoints. The cost of the matching E^* or the earth-mover distance between multisets S and T is $EMD(S,T) = \sum_{e \in E^*} ||e||_1$.

For a grid over $\mathbb{R}^2$, we use a grid's *cell size* to refer to the side length of cells in the grid. Fix a grid G over $\mathbb{R}^2$ whose cell size is a positive integer. For every multiset $S \subseteq [\Delta]^2$, we define $V_G(S)$ to be the *characteristic vector* of S with respect to G. Each coordinate of $V_G(S)$ corresponds to a cell of G that intersects $(0, \Delta)^2$, and the value of that coordinate is the number of points of S in the corresponding cell. In the context of our algorithm, we avoid having points that live on the grid lines so that the number of points that fall into a cell of the grid is defined without ambiguity. Since the points have integral coordinates, we can ensure that the points are in the interior of the grid cells by restricting the grid lines to have half-integral coordinates $\ldots, -\frac{3}{2}, -\frac{1}{2}, \frac{1}{2}, \frac{3}{2}, \frac{5}{2} \ldots$.

Throughout this section, we talk about grids that are shifted by some random 2-dimensional vectors with half-integral coordinates. We assume that each grid prior to shift is fixed at the origin $(0, 0)$. Thus, after shifting a grid by vector $\boldsymbol{v} = (x_0, y_0)$, the grid point at $(0, 0)$ is moved to (x_0, y_0), and the rest of the grid translates accordingly. Thus, we ensure that the lines of the shifted grid have half-integral coordinates.

To estimate the earth-mover distance over data streams, our algorithm maintains sketches of characteristic vectors of the two input sets with respect to different grids. These sketches enable us to estimate the ℓ_1 and ℓ_0 norms of the characteristic vectors [1].

Let V be an N-dimensional vector whose coordinates are values in the set $\{1, \ldots, M\}$. The ℓ_1 norm of $V = (x_1, \ldots, x_N)$ is $||V||_1 = \sum_{i=1}^{N} |x_i|$ and its ℓ_0 norm is $||V||_0 = |\{x_i : x_i \neq 0\}|$. We use the following two facts from [9] and [26] to maintain a sketch for the ℓ_1 and ℓ_0 norm of vector V whose coordinates are dynamically updated in a data stream. Each update in the stream is of the form (i, a) which adds a to the i-th coordinate of V.

Fact 2 (Theorem 2 of [9]). *There is an algorithm that, for any $\epsilon, \delta > 0$, estimates the ℓ_1 norm of V up to a factor of $(1 \pm \epsilon)$ with probability $1 - \delta$ and uses $O(1/\epsilon^2 \cdot \log M \cdot \log(N/\delta) \cdot \log(1/\delta))$ bits of memory.*

Fact 3 (Theorem 1 of [26]). *There is an algorithm that, for any $\epsilon, \delta > 0$, estimates the ℓ_0 norm of V up to a factor of $(1 \pm \epsilon)$ with probability $1 - \delta$ using $O(1/\epsilon^2 \cdot \log N \cdot \log(1/\delta))$ bits of memory.*

Our Technique Our second algorithm is a modification of the idea in [8] which uses an embedding of EMD into ℓ_1 that has a distortion of $O(\log \Delta)$ [12, 6]. However, since any embedding of EMD into ℓ_1 must incur $\Omega(\sqrt{\log \Delta})$ distortion [17], we need additional ideas to obtain a better approximation ratio.

The algorithm of [8] uses nested grids $G_i, i = 0, \ldots, \log \Delta$ over $\mathbb{R}^2$ where the cell size of grid G_i is 2^i, and a cell in G_i contains 4 cells in G_{i-1}. The nested grids are shifted by a vector chosen uniformly at random. The multiset S is mapped into $f(S) = (V_{G_0}(S), 2V_{G_1}(S), \ldots, 2^i V_{G_i}(S), \ldots, 2^{\log \Delta} V_{G_{\log \Delta}}(S))$ in the ℓ_1 space where $V_{G_i}(S)$ denotes the characteristic vector of multiset S with respect to grid G_i. In other words, $f(S)$ is obtained by concatenating vectors $V_{G_0}(S), 2V_{G_1}(S), \ldots, 2^{\log \Delta} V_{G_{\log \Delta}}(S)$. Similarly, multiset T is mapped into $f(T)$, and to estimate $EMD(S, T)$, the value of $||f(S) - f(T)||_1$ is computed. The distortion of the above embedding is $O(\log \Delta)$, so the above algorithm gives a $O(\log \Delta)$-approximation streaming algorithm for computing $EMD(S, T)$.

Instead of using one grid per level, our algorithm uses $O(\log \Delta)$ randomly shifted grids at each level $i = 0, \ldots, \log \Delta$. At each level our algorithm maintains the ℓ_1 norm of the difference of characteristic vectors of S and T with respect to every grid at that level. At the end of the stream, we choose the grid with minimum

[1] The ℓ_0 norm of a vector is also referred to as the frequency moment F_0 in the literature.

ℓ_1 difference at each level and compute our estimate. This way our algorithm circumvents $\Omega(\sqrt{\log \Delta})$ lower bound on the distortion of embedding EMD into ℓ_1 [17]. The proof that the above modification gives a better approximation ratio is the main technical part of this section.

3.1 Algorithm Description

For every $i = 0, \ldots, \log \Delta$, our algorithm builds $2 \log \Delta$ grids over $\mathbb{R}^2$ with cells of size 2^i that are randomly and independently shifted. As the points in the stream arrive, the algorithm maintains a sketch for the ℓ_1 norm of the difference of characteristic vectors of S and T with respect to every grid. At the end of the stream the algorithm chooses, for each level, the grid with minimum ℓ_1 norm and reports $Z = \frac{\widehat{k}^2}{2} \sum_{i=0}^{\log \Delta} 2^i \widehat{C}_i$ as the estimate of $EMD(S,T)$ where $\widehat{C}_i$ is the estimate of the minimum ℓ_1 norm at level i and $\widehat{k}$ is an estimate of the minimum of the number of distinct points in S and T. Thus, our algorithms maintains the following data structures:

1. For each $i = 0, \ldots, \log \Delta$ and each $j = 1, \ldots, 2 \log \Delta$, let G_i^j be a grid of cell size 2^i that is shifted by a vector chosen independently and uniformly at random (recall that the coordinates of the shift vector are half-integral). The algorithm maintains a sketch of vector $V_{G_i^j}(S) - V_{G_i^j}(T)$ under addition and deletion of points from S and T to estimate its ℓ_1 norm, $||V_{G_i^j}(S) - V_{G_i^j}(T)||_1$, at the end of the stream. This can be done using Fact 2.
2. The algorithm also maintains a sketch to determine the number of distinct points in S and T. This can be done using Fact 3 to estimate the ℓ_0 norm of S and T with respect to any grid at level 0. Note that all random shift vectors result in the same grid G_0 at level 0, and there is at most one distinct point of S (or T) in each cell of G_0, so the ℓ_0 norm of $V_{G_0}(S)$ (or $V_{G_0}(T)$) is the number of distinct points of multiset S (or T).

Let $\widehat{k}$ be minimum of the two estimates for the ℓ_0 norms of $V_{G_0}(S)$ and $V_{G_0}(T)$. Then, $\widehat{k}$ estimates k, the minimum of the number of distinct points in S and T. We define $C_i^j = ||V_{G_i^j}(S) - V_{G_i^j}(T)||_1$, and $C_i = \min_j C_i^j$. Let $\widehat{C}_i^j$ denote the algorithm's estimate of C_i^j for all i, j. At every level i the algorithm chooses the grid G_i^j that minimizes $\widehat{C}_i^j$ and lets $\widehat{C}_i = \min_j \widehat{C}_i^j$. The output of the algorithm is

$$Z = \frac{\widehat{k}^2}{2} \sum_{i=0}^{\log \Delta} 2^i \cdot \widehat{C}_i \tag{1}$$

This concludes our description of the algorithm.

Space usage The above algorithm uses $O(\log \Delta)$ grids at each level $i = 0, \ldots, \log \Delta$, and maintains the ℓ_1 norm of the difference of characteristic vectors of S and T with respect to each grid. Each vector $V_{G_i^j}(S) - V_{G_i^j}(T)$ has at most

$O(\Delta^2)$ coordinates and each coordinate is in $\{0, 1, \ldots, n\}$. By Fact 2, the sketch to maintain $C_i^j = ||V_{G_i^j}(S) - V_{G_i^j}(T)||_1$ requires $O(\log n \cdot \log \frac{\Delta}{\delta'} \cdot \log \frac{1}{\delta'})$ bits of storage where δ' is the probability of error in estimating the norm with respect to each grid. If we want the total error probability in estimating all C_i^j to be bounded by δ, we need to set $\delta' = \frac{\delta}{2 \log^2 \Delta}$. With this value of δ' the space needed to maintain each C_i^j is $O(\log n \cdot \log \frac{\Delta}{\delta} \cdot \log \log \Delta \cdot \log \frac{1}{\delta})$ and the total space used to maintain the ℓ_1 norm of these vectors is $O(\log^3 \Delta \cdot \log \log \Delta \cdot \log n)$. Also the space needed to maintain the number of distinct point in S and T is $O(\log \Delta)$ (by Fact 3 from [26]). Thus the total space used by the algorithm is still $O(\log^3 \Delta \cdot \log \log \Delta \cdot \log n)$.

To show that the estimate Z returned by the algorithm approximates $EMD(S, T)$, we will prove upper and lower bounds on the value of Z in the next section.

3.2 Bounding the Cost

In this section we prove upper and lower bounds on the cost of the estimate returned by the algorithm. By Fact 2 and Fact 3, the values of k and $||V_{G_i^j}(S) - V_{G_i^j}(T)||_1$ for all i, j can be estimated within a factor of $1 \pm \epsilon$ for any parameter $\epsilon > 0$. This increases the space usage by a multiplicative factor of $O(1/\epsilon^2)$ which is ignored as we take ϵ to be some small constant. If $\widehat{k} = (1 \pm \epsilon)k$ and $\widehat{C}_i^j = (1 \pm \epsilon)C_i^j$, then:

$$Z = \frac{\widehat{k}^2}{2} \sum_{i=0}^{\log \Delta} 2^i \cdot \widehat{C}_i = (1 \pm \epsilon)^3 \cdot \frac{k^2}{2} \sum_{i=0}^{\log \Delta} 2^i \cdot C_i \tag{2}$$

For fixed $\epsilon > 0$, the factor $(1 \pm \epsilon)^3$ is a small constant. Thus, it suffices to prove upper and lower bounds on the value of $\frac{k^2}{2} \sum_{i=0}^{\log \Delta} 2^i \cdot C_i$ instead of $Z = \frac{\widehat{k}^2}{2} \sum_{i=0}^{\log \Delta} 2^i \cdot \widehat{C}_i$ returned by the algorithm. In fact, to further simplify the exposition, we prove our bounds on the value of $Y = \frac{1}{2} \sum_{i=0}^{\log \Delta} 2^i \cdot C_i$ which is scaled by a factor of $1/k^2$. Specifically we show that $\Omega(1/k^2) \cdot EMD(S, T) \leq Y \leq O(k) \cdot EMD(S, T)$ with very high probability. In the next lemma we show a high probability upper bound on Y.

Lemma 1 (Upper bound). *With high probability, the value* $Y = \frac{1}{2} \sum_{i=0}^{\log \Delta} 2^i \cdot C_i$ *is at most* $O(k) \cdot EMD(S, T)$.

Proof. Recall that E^* denotes the set of edges of the optimal matching between points of S and T. We say an edge $e \in E^*$ *crosses* a grid G if the two endpoints of e fall in different cells of G.

Definition 2 (Good Grid). *A grid G_i^j at level-i is a good grid if it is not crossed by any edge $e \in E^*$ whose ℓ_1 norm is less than $\frac{1}{8k}$-fraction of cell size of G_i^j (i.e. $2^{i-3}/k$).*

To bound Y in terms of $EMD(S,T)$, we show that with very high probability at every level one of the $2\log\Delta$ randomly shifted grids is a good grid. If we consider the set of such good grids, one per level i, then every $e \in E^*$ only crosses grids whose cell size is at most $8k||e||_1$. This allows us to charge the length of each edge $e \in E^*$ to the grids that it crosses at different levels.

Let E_i^j be the event that G_i^j (i.e. the j-th randomly shifted grid at level i) is a good grid. The following claim states that with very high probability, there is a good grid G_i^j at every level i.

Claim 1. $\Pr[\ \forall i\ \exists j\ E_i^j\ ] \geq 1 - \dfrac{\log\Delta}{\Delta^2}$.

Proof. We assume without loss of generality that k is the number of distinct points in T. Each edge in the optimal matching connects one of these k points to a point in S. Any edge $e \in E^*$ where $||e||_1 < 2^{i-3}/k$ connects a point $p \in T$ to a point in S which is in a square of side length $2^{i-2}/k$ centered at p. Grid G_i^j is shifted by a random vector, and it intersects one of the edges whose ℓ_1 norm is $< 2^{i-3}/k$ only if it intersects one of the k squares of side length $2^{i-2}/k$ centered at points in T. The cell size of grid G_i^j is 2^i, and the side length of each square is $2^{i-2}/k$, so the probability that a square is intersected by a line of grid G_i^j is $\leq 1/2k$. By union bound the probability that any of the k squares intersect a line of grid G_i^j is at most $1/2$. This also bounds the probability that grid G_i^j is crossed by an edge of length $< 2^{i-3}/k$. Thus, the probability over random shift vectors that grid G_i^j is not a good grid is at most $1/2$. There are $2\log\Delta$ shift vectors at level i, and by independence of shift vectors the probability that all grids G_i^j at level i are not good is at most $\left(\frac{1}{2}\right)^{2\log\Delta} = \frac{1}{\Delta^2}$. The claim is proved by applying the union bound for all $\log\Delta$ levels.

We next show how to bound Y from above using Claim 1. Let's assume that there is a good grid at each level i denoted by G_i^*, then:

$$Y = \frac{1}{2}\sum_{i=0}^{\log\Delta} C_i \cdot 2^i \leq \frac{1}{2}\sum_{i=0}^{\log\Delta} ||V_{G_i^*}(S) - V_{G_i^*}(T)||_1 \cdot 2^i \tag{3}$$

It's easy to see that for every grid G: $||V_G(S) - V_G(T)||_1 \leq 2|\{e \in E^* :$ e crosses $G\}|$ for the following reason. For every square $c \in G$, let $n_c(S)$ and $n_c(T)$ be the the number of points of multiset S and T in square c respectively. Then in every cell $c \in G$, $|n_c(S) - n_c(T)|$ is the minimum number of points in cell c that cannot be matched with a point in that square. Thus, in any matching the total number of points that are not matched within their square is at least $\sum_{c\in G}|n_c(S) - n_c(T)| = ||V_G(S) - V_G(T)||_1$. Each point that is not matched within its square is an endpoint of an edge that crosses grid G. Thus the number of edges in any matching that cross G is at least $||V_G(S) - V_G(T)||_1/2$. This combined with (3) implies that:

$$Y \leq \sum_{i=0}^{\log\Delta} \left|\{e \in E^* : e \text{ crosses } G_i^*\}\right| \cdot 2^i \tag{4}$$

Since G_i^* is a good grid, the above implies that:

$$\begin{aligned}
Y &\le \sum_{i=0}^{\log \Delta} \left|\{e \in E^* : e \text{ crosses } G_i^*\}\right| \cdot 2^i \le \sum_{i=0}^{\log \Delta} \left|\{e \in E^* : ||e||_1 > \frac{1}{k} 2^{i-3}\}\right| \cdot 2^i \\
&\le \sum_{i=0}^{\log \Delta} 2^i \sum_{e \in E^* : ||e||_1 > 2^{i-3}/k} 1 \le \sum_{e \in E^*} \sum_{i : 2^{i-3} < k||e||_1} 2^i = \sum_{e \in E^*} \sum_{i=0}^{3+\log k||e||_1} 2^i \\
&\le \sum_{e \in E^*} 2^{4+\log k||e||_1} = \sum_{e \in E^*} 16k||e||_1 \le 16k \cdot EMD(S,T)
\end{aligned}$$

Thus, $Y \le 16k \cdot EMD(S,T)$ if there is a good grid at each level i which happens with probability $1 - \frac{\log \Delta}{\Delta^2}$ by Claim 1. Hence, $\Pr[Y \le 16k \cdot EMD(S,T)] \ge 1 - o(1)$.

Using the above result, we can also show an upper bound on the expected value of Y, but we omit the straightforward details. Our next lemma establishes the lower bound on the value returned by the algorithm.

Lemma 2 (Lower bound). *The value* $Y = \frac{1}{2} \sum_{i=0}^{\log \Delta} 2^i \cdot C_i$ *is at least* $\Omega(\frac{1}{k^2}) \cdot EMD(S,T)$.

Proof. The idea of the lower bound is to charge the cost of each edge $e \in E^*$ to the grid levels whose cell size is at most $||e||_1$. Thus at each level only edges whose ℓ_1 norm is at least the cell size of that level contribute to the cost. Then, at each level i we bound from above the total number of edges that contribute to that level in terms of $C_{i'} = \min_j ||V_{G_{i'}^j}(S) - V_{G_{i'}^j}(T)||_1$ where i' is *a few* levels below i. Therefore, we can bound $EMD(S,T)$ from above in terms of Y.

For any i, we use E_i^* to denote $\{e \in E^* : ||e||_1 \ge 2^i\}$. Note that $|E_i^*| \le n$ for all i because there are a total of n edges in E^*. We have:

$$\begin{aligned}
EMD(S,T) &= \sum_{e \in E^*} ||e||_1 \le \sum_{e \in E^*} \sum_{i=0}^{\log ||e||_1} 2^i \le \sum_{i=0}^{\log 2\Delta} \sum_{e \in E^* : ||e||_1 \ge 2^i} 2^i \\
&= \sum_{i=0}^{\log 2\Delta} \left|E_i^*\right| \cdot 2^i \le \left(\sum_{i=0}^{\log k} n \cdot 2^i + \sum_{i=\log 2k}^{\log 2\Delta} \left|E_i^*\right| \cdot 2^i \right) \\
&\le 2nk + \sum_{i=\log 2k}^{\log 2\Delta} \left|E_i^*\right| \cdot 2^i = 2nk + \sum_{i=0}^{\log 2\Delta - \log 2k} \left|E_{i+\log 2k}^*\right| \cdot 2^{i+\log 2k} \\
&= 2nk + 2k \sum_{i=0}^{\log \Delta - \log k} \left|\{e \in E^* : ||e||_1 \ge 2k \cdot 2^i\}\right| \cdot 2^i. \qquad (5)
\end{aligned}$$

The main tool in the proof of the lemma is the following claim which lower bounds $C_i = \min_j ||V_{G_i^j}(S) - V_{G_i^j}(T)||_1$ by the number of edges in the optimal matching E^* whose length is at least $k2^{i+1}$.

Claim 2. *For all* $i = 0, \ldots, \log \Delta$: $\left|\{e \in E^* : ||e||_1 \geq k \cdot 2^{i+1}\}\right| \leq \frac{k}{2} \cdot C_i$.

The idea of the proof is to view grid G at level i as a graph where the grid cells are vertices of the graph and the edges crossing the grid are directed edges of the graph. We then show how to decompose the edges of the graph into a set of paths of length $\leq k$ where the start and end vertex of each path contribute two to the value of C_i. Thus the total number of such paths is at most $C_i/2$ and the total number of edges whose ℓ_1 norm is $\geq k \cdot 2^{i+1}$ is at most $kC_i/2$. The detailed proof appears in the appendix.

We next show how to use the above claim to prove Lemma 2. From Inequality (5) above we have:

$$\begin{aligned} EMD(S,T) &\leq 2nk + 2k \sum_{i=0}^{\log \Delta - \log k} \left|\{e \in E^* : ||e||_1 \geq k \cdot 2^{i+1}\}\right| \cdot 2^i \\ &\leq 2nk + 2k^2 \cdot \frac{1}{2} \sum_{i=0}^{\log \Delta - \log k} C_i 2^i \qquad \text{(by Claim 2)} \\ &= O(k^2) \cdot Y \end{aligned}$$

The last inequality holds because $n = C_0/2 \leq Y$. This completes the proof of Lemma 2.

Lemma 1 and 2 together imply that our algorithm gives a $O(k^3)$-approximation of the cost of EMD. To ensure approximation ratio of $O(\min(k^3, \log \Delta))$, the algorithm holds an additional data structure to maintain the sketch used by $O(\log \Delta)$-approximation algorithm of [8]. The two algorithms maintain their own sketches and at the end of the stream, each algorithm computes its estimate of EMD using its sketch and the minimum of the two estimates is returned. Clearly this estimate is within $O(\min(k^3, \log \Delta))$ factor of the cost of EMD. Thus, we have the following:

Theorem 2. *There is a* $O(\min(k^3, \log \Delta))$*-approximation that uses* $O(\log^{O(1)} \Delta n)$ *space to estimate the earth mover distance between two multiset* $S, T \subseteq [\Delta]^2$ *given over a dynamic data stream, where* $|S| = |T| = n$ *and minimum of the number of distinct points in* S *and number of distinct points in* T *is bounded by* k*.*

4 Conclusion

We have obtained two approximation algorithm for earth-mover distance between two multisets of points in $[\Delta]^2$ when the number of distinct points in one set is small. Both algorithms use polylogarithmic space. Our algorithms can be extended to give streaming algorithms for capacitated k-median clustering with hard constraints. We conclude with some natural open questions: 1) Is there a $O(1)$-approximation algorithm for EMD with no constraints on the input size using only polylogarithmic space? 2) Can one prove a lower bound on the best approximation possible for EMD in polylogarithmic space? 3) Are there better streaming algorithms for the capacitated k-median with hard constraints?

References

[1] R. Krauthgamer A. Andoni, P. Indyk. Earth mover distance over high-dimensional spaces. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete Algorithms*, pages 343–352, 2008.

[2] C. Harrelson R. Krauthgamer K. Talwar A. Archer, J. Fakcharoenphol and E. Tardos. Approximate classification via earthmover metrics. In *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1079–1087, 2004.

[3] D. Stubbs A. McGregor. Sketching earth-mover distance on graph metrics. *Lecture notes in computer science*, pages 274–286, 2013.

[4] A. Andoni, Khanh Do Ba, P. Indyk, and D. Woodruff. Efficient sketches for earth-mover distance, with applications. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science*, 2009.

[5] J. Naor C. Chekuri, S. Khanna and L. Zosin. Approximation algorithms for the metric labeling problem via a new linear programming formulation. In *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 109–118, 2001.

[6] M. Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 380–388, 2002.

[7] C. Sohler G. Frahling. Coresets in dynamic geometric data streams. In *The 37th Annual ACM symposium on Theory of Computing*, pages 209–217, 2005.

[8] P. Indyk. Algorithms for dynamic geometric problems over data streams. pages 373–380, 2004.

[9] P. Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream computation. *Journal of the ACM (JACM)*, 53(3):307–323, 2006.

[10] P. Indyk. A near linear time constant factor approximation for euclidean bichromatic matching (cost). In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete Algorithms*, pages 39–42, 2007.

[11] P. Indyk, A. McGregor, I. Newman, and K. Onak. Open problems in data streams, property testing, and related topics. *Bernitoro Workshop on Sublinear Algorithms*, 2011. Available at http://people.cs.umass.edu/ mcgregor/papers/11-openproblems.pdf.

[12] P. Indyk and N. Thaper. Fast image retrieval via. embeddings. *Workshop on Statistical and Computational Theories of Vision at ICCV*, 2003.

[13] Y. Rabani J. Chuzhoy. Approximating k-median with non-uniform capacities. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete Algorithms*, pages 952–958, 2005.

[14] E. Lawler. *Combinatorial Optimization: Networks and Matroids*. 1976.

[15] T. Wexler M. Pal, E. Tardos. Facility location with non-uniform hard capacities. In *Proceedings of the 42nd IEEE symposium on Foundations of Computer Science*, page 329, 2001.

[16] A. McGregor. Open problems in data streams and related topics. *IITK Workshop on Algorithms For Data Streams*, 2006. Available at http://www.cse.iitk.ac.in/users/sganguly/workshop.html.

[17] A. Naor and G. Schechtman. Planar earthmover is not in l_1. *SIAM Journal on Computing*, 37(3):804–826, 2007.

[18] K. Varadarajan P. Agarwal. A near-linear constant-factor approximation for euclidean bipartite matching?

[19] R. Sharathkumar P. K. Agarwal. A near-linear time ϵ-approximation algorithm for geometric bipartite matching. In *The 44th Annual ACM Symposium on Theory of Computing*, pages 385–494, 2012.

[20] S. Mazumdar S. Har-Peled. On coresets for k-means and k-median clustering. In *The 36th Annual ACM Symposium on Theory of Computing*, pages 291–300, 2004.

[21] A. Naor S. Khot. Nonembeddability theorems via fourier analysis. *Mathematische Annalen*, 334:821–852, 2006.

[22] H. Rom S. Peleg, M. Werman. A unified approach to the change of resolution: Space and gray-level. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:739–742, 1989.

[23] P. Vaidya. Geometry helps in matching. *SIAM Journal of Computing*, 18:1201–1225, 1989.

[24] L. J. Guibas Y. Rubner, C. Tomasi. A metric for distributions with applications to image data-bases. In *Proceedings of the Sixth International Conference on Computer Vision*, pages 59–66, 1998.

[25] L. J. Guibas Y. Rubner, C. Tomasi. The earth mover's distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99–121, 2000.

[26] T. S. Jayram Z. Bar-Yossef, R. Kumar, D. Sivakumar, and L. Trevisan. Counting distinct elements in a data stream. In *Proceedings of the 6th International Workshop on Randomization and Approximation Techniques (RANDOM'02)*, pages 1–10, 2002.

A Proof of Claim 2

Proof. We prove that for *any* grid G_i^j at level i, the number of edges of E^* with ℓ_1 norm at least $k2^{i+1}$ is at most $kC_i^j/2$, where $C_i^j = ||V_{G_i^j}(S) - V_{G_i^j}(T)||_1$.

For grid G_i^j, we define the directed multigraph Γ as follows. Every cell c of the grid that contains a point from S or T corresponds to a vertex of Γ. Each edge of Γ corresponds to an edge of the matching E^* that crosses the grid. Recall that an edge crosses the grid if its endpoints are in different cells. Each edge in Γ has a *length* which is equal to the ℓ_1 norm of the corresponding edge in the matching E^*. Edges of Γ are directed from points in S to points in T. Note that S and T are multisets, and there might be multiple copies of a point in S or in T. Thus, if p copies of a point in S are matched to p copies of a point in T, there are p copies of the same edge in the matching E^*, and if the endpoints of the edge are in different cells, there are p edges between the corresponding vertices in Γ. Hence, Γ is in general a multigraph.

A few observations about graph Γ are in order:

Observation 1. *The length of every simple cycle (or every simple path) in graph Γ is at most k.*

This holds because k is an upper bound on the number distinct points in T and also the number of vertices of Γ with positive indegree.

Observation 2. *No cycle in Γ contains an edge of length $\geq k \cdot 2^{i+1}$.*

If such a cycle exists, there are matched pairs $(s_1, t_1), (s_2, t_2), \ldots, (s_\kappa, t_\kappa) \in E^*$ such that for every $r \in [\kappa]$, points t_r and s_{r+1} (we define $\kappa + 1 = 1$) are in the same cell in grid G_i^j and the length of one of the edges is $\geq k2^{i+1}$ (see Figure 1). Thus, the matching E^* costs at least $k2^{i+1}$.

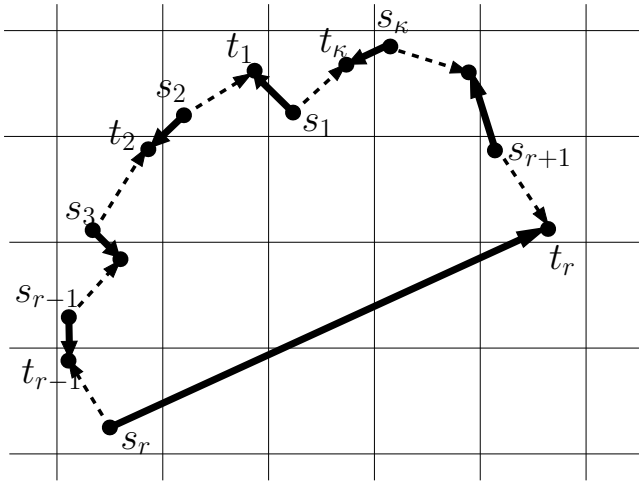

Fig. 1. The solid edges show the matching of the pairs $(s_1, t_1), (s_2, t_2), \ldots$. These edges form a cycle in the corresponding graph Γ. If there is an edge (s_r, t_r) with length $\geq k2^{i+1}$, there is an alternative matching $(t_1, s_2), (t_2, s_3), \ldots$ (shown in dotted lines) with lesser cost.

However, there is an alternative matching $(t_1, s_2), (t_2, s_3), \ldots, (t_\kappa, s_1)$ that costs at most $k2^{i+1}$ because the are at most k pairs of points (by Observations 1), and each pair is in the same cell of grid G_i^j. This contradicts the optimality of E^*.

The next observation is a connection between the number of points of S and T in a cell of the grid, and the indegree and outdegree of the corresponding cell in graph Γ. Let c be any cell in G_i^j and v_c be the corresponding vertex in graph Γ. Let $\deg^+(v_c)$ and $\deg^-(v_c)$ denote the outdegree and indegree of v_c respectively, and let $n_c(S)$ and $n_c(T)$ denote the number of points of S and T in the cell c.

Observation 3. *For every cell $c \in G_i^j$ and corresponding vertex v_c in Γ:*
$\deg^+(v_c) - \deg^-(v_c) = n_c(S) - n_c(T)$.

For every edge out of v_c, there should be a point of S in cell c, and for every edge into v_c there should be a point of T in cell c. The rest of the points of S and T in cell c that are not an endpoint of an edge of Γ are matched within the cell. The number of points of S that are matched within c should be equal to the number of points of T that are matched within c, so these points don't contribute any value to $n_c(S) - n_c(T)$. Hence, $\deg^+(v_c) - \deg^-(v_c) = n_c(S) - n_c(T)$.

By summing over all cells c in the grid, this observation implies that:

$$\sum_{v \in V(\Gamma)} |\deg^+(v) - \deg^-(v)| = \sum_{c \in G_i^j} |n_c(S) - n_c(T)| = ||V_{G_i^j}(S) - V_{G_i^j}(T)||_1 = C_i^j \tag{6}$$

We are now ready to prove claim 2. The idea is to decompose the edges of graph Γ into a set of paths where each path contains at least one of the edges of length $\geq k2^{i+1}$. We show that the graph Γ can be decomposed into at most $C_i^j/2$ such paths and each path contains at most k edges of length $\geq k2^{i+1}$.

The decomposition works in a natural way as follows. Let $e = (u, t)$ be any edge such that $||e|| > k2^{i+1}$. We show how to construct a path that contains e. If $\deg^-(t) > \deg^+(t)$, t is the end of the path. Otherwise, $\deg^-(t) \leq \deg^+(t)$ and there is an edge (t, w) going out of t. By the same argument either $\deg^-(w) > \deg^+(w)$ or there is an edge going out of w. This process can be repeated until a vertex z is reached such that $\deg^-(z) > \deg^+(z)$. We mark z to be the end of the path. The original edge $e = (u, t)$ can also be traced in the opposite direction to reach a vertex a such that $\deg^+(a) > \deg^-(a)$. Note that by Observation 2 edge e is not in any cycle, so the start and end vertex of the path can't be the same vertex. Removing all the edges of this path from graph Γ reduces $\sum_{v \in V(\Gamma)} |\deg^+(v) - \deg^-(v)|$ by two because the quantity $|\deg^+(v) - \deg^-(v)|$ is reduced by one for the start and end vertices of the path, and for all other vertices this quantity is unchanged.

After removing the edges of this path from graph Γ, the remaining graph may still contain edges of length $\geq k2^{i+1}$. We can choose any one of these edges and repeat the above process to find another path and remove it from the graph. This process can be repeated until there are no edges of length $\geq k \cdot 2^{i+1}$ in the graph. Each time a path is extracted the quantity $\sum_{v \in V(\Gamma)} |\deg^+(v) - \deg^-(v)|$ reduces by two. Thus, the total number of such paths is at most $\frac{1}{2} \sum_{v \in V(\Gamma)} |\deg^+(v) - \deg^-(v)|$ which equals $C_i^j/2$ by Equation (6). Each path contains at most k edges of length $\geq k \cdot 2^{i+1}$ by Observation 1. Thus the total number of such edges is bounded by $kC_i^j/2$.